



Hewlett Packard
Enterprise

HPE Operations Orchestration

Software Version: 10.60

Windows and Linux Operating Systems

Architecture Guide

Document Release Date: May 2016

Software Release Date: May 2016

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2016 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation Updates

To download the most recent edition of a document, go to <https://softwaresupport.hp.com>.

Contents

System Architecture	4
Operations Orchestration Components	4
Simple Deployment	4
Simple Cluster	4
Scalability	5
Adding a RAS	6
RAS High Availability	7
Using a Load Balancer in OO Deployment	8
Load Balancer Requirements	8
Load Balancer Security	8
Configuring the Load Balancer and OO Centrals for TLS Offloading	9

System Architecture

Operations Orchestration Components

HPE OO Studio is a standalone authoring program used for creating, modifying, and testing flows.


HPE OO Central is the run time environment of HPE OO. It is used for running flows, monitoring the various runs, and generating reports.

A **RAS** is a remote action server, containing a worker and a remote protocol for connecting with Central.

For additional information on HPE OO components, see the *HPE OO Concepts Guide*.


Simple Deployment

The basic HPE OO deployment consist of a single Central instance, as shown in the image below.



Simple Cluster

In order to prevent the Central being the single point of failure, it is recommended to have a high-availability deployment. You can set a cluster of multiple Central nodes, the simplest of which contains two Central nodes connected to the same database schema. As shown in the image below, a load balancer can be set before the Central cluster to expose a single URL to the end users. Exposing a single URL can also be done with DNS load balancing.



The load balancer/reverse proxy should redirect to the Centrals that use ports 8443 and 8080, if the default values were chosen during installation. For more information, see the *HPE OO System Requirements*.

Change from OO 9.x: Unlike in previous versions, there is no need for external clustering software, nor is there a requirement for a shared file system.

Scalability

HPE OO offers horizontal scaling for increasing execution throughput.

You can add more Central instances to the HPE OO cluster. HPE OO supports live scalability, which means that no downtime is required when adding a Central node. Simply install an additional Central instance and point it to the existing database schema.

For more information, see the *HPE OO 10 Benchmark* document, available on HPLN at <https://hpln.hp.com/node/17617/attachment>.

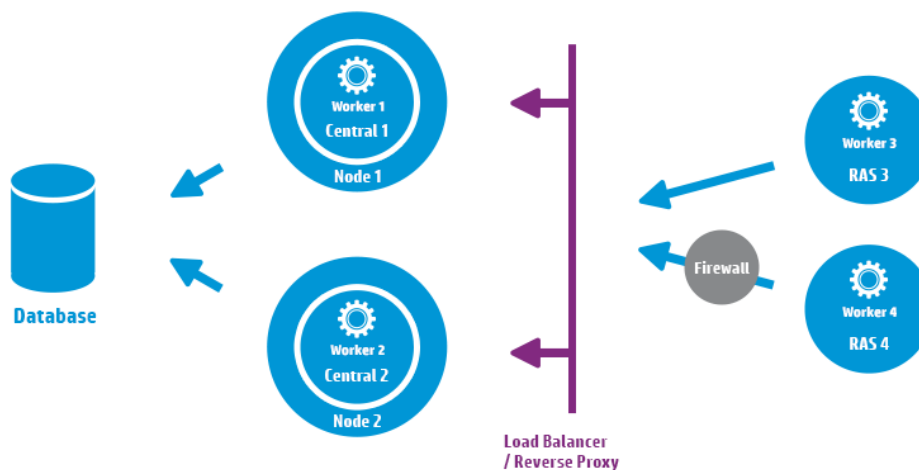
Adding a RAS

A RAS instance is an optional OO component. A RAS can be used if HPE OO needs to run flows in a network segment that is not reachable from the HPE OO Central nodes. In such case, you can install a RAS instance in the target network segment and it will pull the required flows from the Central and run them locally.

Another use case where a RAS can be used is when the executed flow requires specific binaries on the local machine. There is no need to install the binaries on each HPE OO node. It is enough to install them on a host where a RAS is installed, and configure the flows (or specific steps) to run on this RAS. This can be achieved by leveraging the worker group functionality.

For more information on worker groups, see the *HPE OO Concepts Guide*.

You can attach RAS instances to HPE OO Central or a cluster of HPE OO Central nodes. The image below shows how RAS3 and RAS4 communicate with the HPE OO Central cluster. Note that RAS4 is located behind a firewall.



Configuring the RAS connectivity direction

In OO 10.60 and later, you can configure RASes so that some initiate the connection to Central while others wait for Central to initiate the connection.

For example, if Central and a RAS are installed in different networks, with Central deployed on a more secured network, and your security rules do not allow connecting from the less secured network to the more secured one, you can have Central initiate the connection to the RAS.

During the installation of a RAS, you must choose between two options:

- **Standard RAS - RAS initiates communication to Central** - this is the simplest option and is recommended if your security rules permit it.
- **Reverse RAS - Central initiates communication to RAS** - choose this option if Central is installed in a different, more secured network, and your security rules do not allow connecting from the less secured

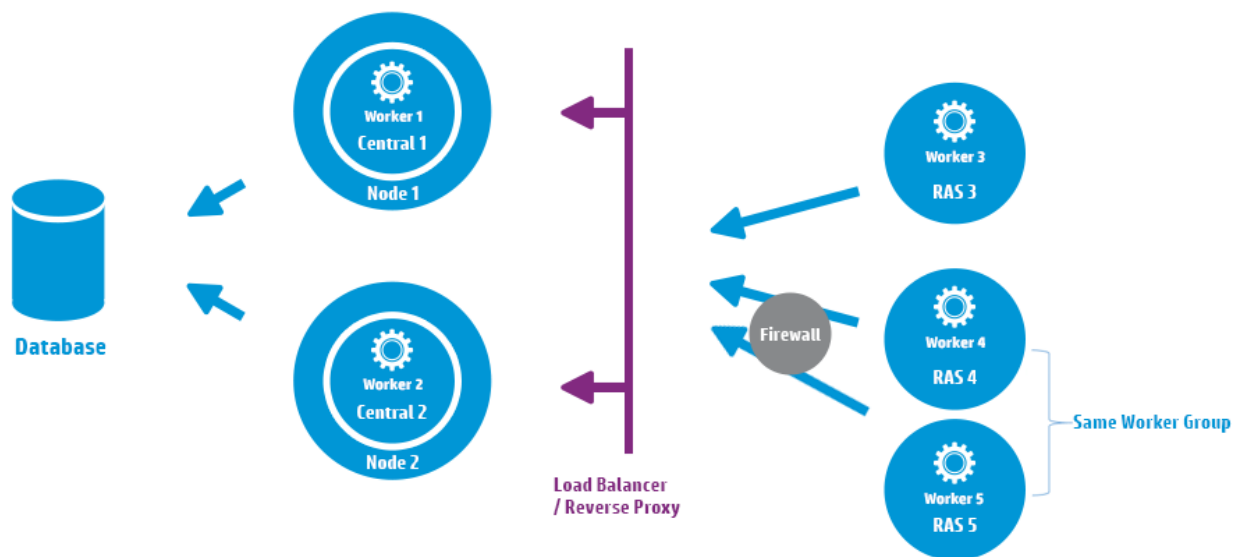
network to the more secured one.

You will need to configure the RAS to accept connection from Central, and configure Central to register the RAS (in the **System Configuration > Topology > Workers** tab).

When the RAS starts up, it will be idle, waiting for Central to initiate connection.

RAS High Availability

When a RAS is deployed in a network segment to manage the machines in that segment, you do not have to make do with a single instance. To achieve high availability, you can deploy an additional RAS instance in the same segment. Make sure to associate it with the same worker group. This is illustrated in the image below:



Change from HPE OO 9.x: There is no need for an additional load balancer between the RAS cluster and Central (or central cluster). Because both RAS 4 and RAS 5 belong to the same worker group, they share the load of executing flows/steps that are designated for that worker group and provide high availability.

Using a Load Balancer in OO Deployment

For information about how to install a load balancer, see the documentation provided by your load balancer vendor.

Load Balancer Requirements

We recommend to configure the load balancer with two separate virtual IPs for the user interface and for RASes:

- For the HPE OO user interface and customer portals, the virtual IP should use a **sticky session** policy. The sticky session ensures that all subsequent requests will be sent to the server that handled the first login request. This means that users will only need to log in to the HPE OO interface once.
- For RASes, the virtual IP should use a **round robin** policy, to distribute the load across the different servers.

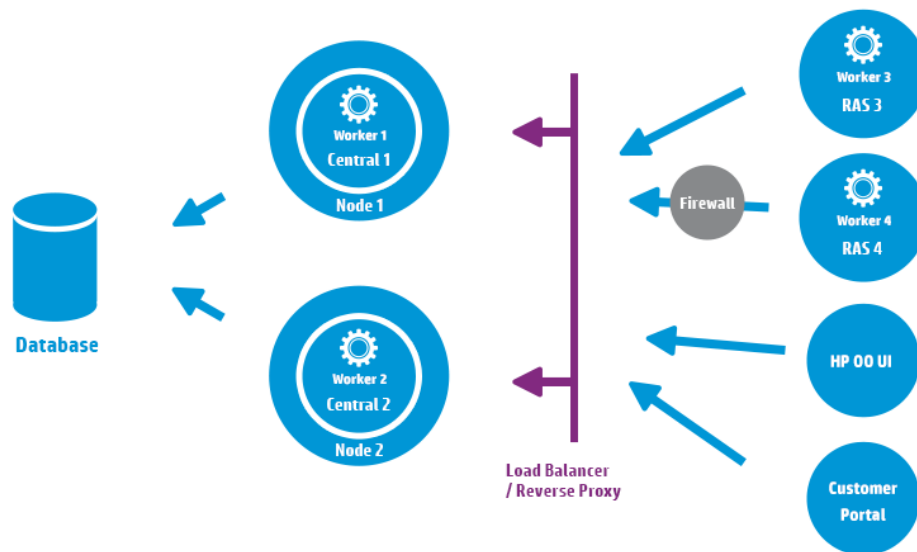
Note: If you have a different configuration that satisfies these requirements, it is okay to use it. For example, if you have a load balancer that supports JSESSION, you can use the JSESSIONID parameter to set up a single virtual IP with a sticky session policy for all sources. Since RAS requests are stateless (no JSESSIONID), this will provide a round robin policy for RASes.

Central uses the following URL to check which server is live: HTTP://<IP>:<PORT>/oo/hello.html

Load Balancer Security

In a hardened high availability environment, the load balancer should be configured for TLS. For information about how to configure TLS, see "Server and Client Certificate Authentication" in the *HPE OO Security and Hardening Guide*.

Communication between the HPE OO interface and the load balancer can use HTTPS. We recommend to install the TLS certificate on the load balancer so that this is the termination point for the encryption. Beyond the load balancer, communication will continue in HTTP, at a faster rate.



Configuring the Load Balancer and OO Centrals for TLS Offloading

If a load balancer is used to access the Central servers, it is recommended to configure the load balancer for TLS offloading.

1. Edit the Tomcat **server.xml** file, to include the following, for example:

```
<Engine name="Catalina" defaultHost= "localhost" >
. . .
<Valve className="org.apache.catalina.valves.RemoteIpValve" protocolHeader="X-
Forwarded-Proto" />
. . .
</Engine>
```

2. Configure the load balancer to add a new header to all the clients' requests.

The header name is configurable and should match the Tomcat configuration specified above. In this example, the name is "X-Forwarded-Proto".

In the F5 load balancer, the configuration would look like this:

```
when HTTP_REQUEST {
  HTTP::header insert "X-Forwarded-Proto" "https";
}
```

